

A2
D3

--For these experiments, the ASR engine's first-pass Viterbi search graph was biased with word bigram data extracted from subsets of the recognition results on given development input data. The result transcriptions (i.e., recognition hypotheses) were randomized and then various sized portions, starting from the top-choice down, were taken to accumulate word pair frequencies. A variation on this experiment imposed a score threshold on the recognition results as the sub-setting mechanism. These counts were converted into probabilities, and these probabilities were used to bias the Viterbi search in favor of the most likely word sequences. The main results of these experiments are summarized in the table illustrated below

[Insert tables from Figures 3A and 3B]

--

the paragraph beginning on page 17, line 12:

D4
A3

--The tables illustrated below show the experimental results for the Names
task.

[Insert tables from Figures 4A and 4B]